

# Predicting Online Video Engagement Using Clickstreams

Everaldo Aguiar  
University of Notre Dame  
Notre Dame, Indiana 46556  
eaguiar@nd.edu

Saurabh Nagrecha  
University of Notre Dame  
Notre Dame, Indiana 46556  
snagrech@nd.edu

Nitesh V. Chawla\*  
University of Notre Dame  
Notre Dame, Indiana 46556  
nchawla@nd.edu

## ABSTRACT

In the nascent days of e-content delivery, having a superior product was enough to give companies an edge against the competition. With today's fiercely competitive market, one needs to be multiple steps ahead, especially when it comes to understanding consumers. Focusing on a large set of web portals owned and managed by a private communications company, we propose methods by which these sites' clickstream data can be used to provide a deep understanding of their visitors, as well as their interests and preferences. We further expand the use of this data to show that it can be effectively used to predict user engagement to video streams.

## Author Keywords

Clickstream; predictive analysis; online video; user engagement.

## ACM Classification Keywords

I.5.2. Pattern Recognition: Design Methodology

## INTRODUCTION

The constant growth in volume, speed, availability, and functionality of the Web brings with it not only a variety of challenges and risks, but also a number of opportunities. While there have been a series of major advances in the field over time, one that has been given a considerable amount of attention in more recent years is that of *personalization*.

Data about users' online activity is continuously captured and analyzed. Advanced recommendation systems are now able to tell us what products we might be interested in buying, the books we will enjoy reading, what movies we should watch next, and even which diseases we are at risk of contracting. From a business perspective, the benefits of being able to understand customers in this level of detail are unquestionable.

Methods for capturing user data on the Web are also becoming increasingly efficient. As described in [20], the browsing behavior of individual users can be recorded at the granularity of mouse clicks with little to no work needing to be done. A number of services, both free and proprietary, offer user tracking solutions that can be implemented and deployed within minutes. However, the feedback that one usually gets from these tools is often in the form of simplistic aggregate statistics that do not offer a deeper understanding of user behavior.

With that in mind, we set to analyze the application of some of these ideas to a specific context, while having as our major

goal the understanding of each user as an individual unit. For this study, we were provided a large dataset that describes user clicks generated within a two-month span and across a number of websites managed by a large communications company.

This paper describes the process through which we parsed, analyzed, and drew knowledge from a user-generated clickstream dataset provided by a large communications company. We begin by showing, from a more general perspective, how this type of data can be used to identify particularly interesting trends in user interest, and to further illustrate the usefulness of this information, we describe how we applied methods to predict user engagement to video streams and discuss their accuracy.




Figure 1. Video viewership drop-off by category of content

Distributing content that entices user engagement and captures large audiences is the ultimate goal of all web media providers. Measuring and forecasting these variables, however, is not an easy task. As Figure 1 illustrates, as time goes by, the amount of users that remain tuned to video streams dramatically decreases. For certain categories, the percentage of users that actually watch videos to completion can be as low as 20%.

To address this undesired outcome, we propose the development of clickstream-based models that can learn the individual preferences and characteristics of each user, and utilize this information to predict how "engaged" they will be to a particular video stream. Being able to know, in advance, if

\*Corresponding Author

a user is likely to exit a video prematurely allows content providers some leeway to implement personalized intervention strategies aimed at maximizing viewership retention.

The remaining portion of the paper is organized as follows – The next section gives an overview of the most recent related literature. That is followed by a detailed coverage of clickstream data representation and a description of our particular dataset. We then elaborate on the methods applied in this study, the results obtained, and their importance. Finally, the last section draws conclusions about this exercise and argues for the latent potential that resides in user-generated clickstream data.

## RELATED WORK

Interest in analyzing the online activities of users is as old as providing consumable content itself. This problem has piqued the interest of multiple fields, namely marketing, psychology and computer science.

Since user activity provides an immense amount of measurable secondary data, various models to predict multiple aspects of their behavior have been proposed. User interaction has been studied at various levels— from gaze tracking [8] to broader patterns of path traversal within a website [1, 22]. Simple duration and dwell-time [4] can be used to predict when a user exits the site. User classification [21] can be used to identify what the user is specifically looking for and even morph the website [16] according to the custom tastes of that particular user profile. Personalized content based on click history has been implemented and widely adopted by commercial content providers [6, 19].

With the distribution of video content online becoming mainstream, the way we study user engagement has been greatly enriched. Studies like [7] have measured the role of video content quality in influencing user engagement, but did not utilize clickstreams to contextualize the video views. Online video engagement for Massive Open Online Courses (MOOCs) [13] has shown that the lessons learned from analyzing video views can be used to improve video authoring, editing and interface design. It also emphasizes the value of *video dropout* as a metric for engagement. Though the MOOC work lacks the contextual history of the users, in this paper we leverage similar and many other clickstream features to predict video engagement.

## CLICKSTREAM DATA REPRESENTATION

Clickstream data consists of a “virtual trail” that users leave behind while they interact with a given system, website or application. More specifically, data that describes the state of a user’s current session is recorded each time a click is performed, and the aggregation of that produces a *clickstream*, which can be used to reconstruct all actions taken by the user while he or she utilized that given product.

While applicable to a variety of scenarios, the collection and analysis of clickstreams has become most notably popular in the context of Web-based tools and websites. As highlighted by Srivastava et. al. [25], the analysis of such information

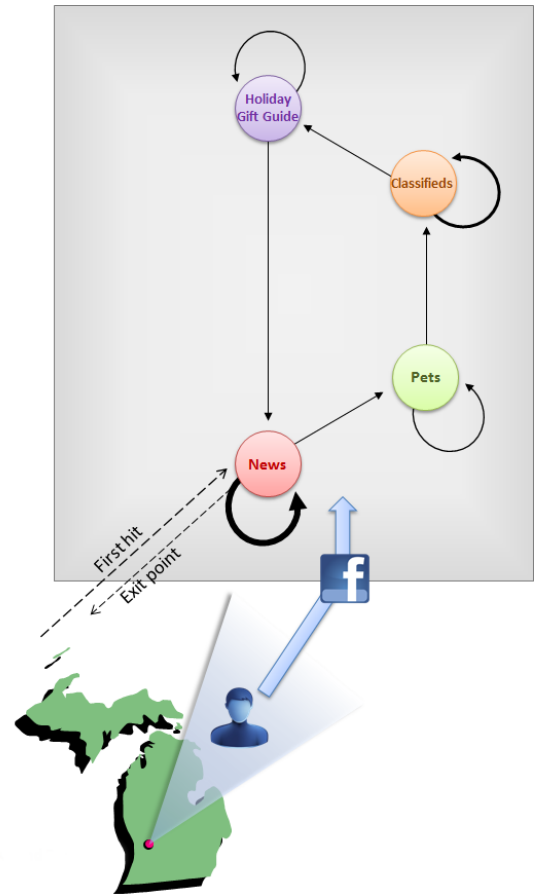


Figure 2. A simple illustration of the clickstream of a typical user

has potential applications in a number of areas such as website personalization and modification, system improvement, business intelligence and usage characterization. Our contributions fall mainly within the first and last domains.

## Our Dataset

The data we utilized for this study was provided to us by a large U.S.-based communications company that operates in the radio, TV, newspaper and online media domain. They manage a few dozens of websites, all of which are embedded with clickstream capturing functionality. Next, we give a detailed description of the most important features this dataset contains.

User activity is continuously captured by numerous servers across the country and is then concatenated at the end of the day in the form of daily “dumps”. We utilized 59 of these files that covered the period ranging from December 4, 2012 to January 31, 2013. Altogether, these files contain an upwards of 65 million click instances.

Each click instance recorded is characterized by a large number of *features* (161 in this case). Table 1 lists a small subset of the most relevant features and a brief description of each. With that information we are able to determine (1) how users reached the website, (2) what attracted them there, (3) what

Feature Type	Feature Name	Description
<i>Nominal</i>		
	Browser	The browser that was used
	Channel	The site that the page view belongs to
	City	The city the user accessed the page from
	Cookies	Whether the user had cookies turned on or not
	Country	The country the user accessed the page from
	Domain	Domain of the user's ISP
	Exclude hit	Identifies web crawlers
	First hit page	URL the user first landed on the website
	Frequency of visits	Denotes hourly, daily, weekly, monthly or yearly visit
	IP	Refers to the IP address of the user
	New visit	Determines whether the user is new to the site, based on cookies
	Referrer	Lists the URL of the website that referred this user
	Region	Refers to the state or region the user was in
	Search Keywords	The search string which led to the particular page
	Section	The section of the website where the click took place
	Subsection	Subsection of the website where the click took place
<i>Numeric</i>		
	First hit time	Timestamp of when the user first landed on the website
	Last click	Time stamp of when the last click was made by the user
	Last visit	Refers to when the user visited the site last
	Time & Date	Timestamp of when the click instance happened
	Visit number	Refers to the number of times the user has visited the site

**Table 1. Dataset features described**

actions they performed while on the site and (4) how they eventually exited.

Note that while there is no feature that captures the event of a user leaving the website, as is common practice, we work under the assumption that when a user is inactive for a period longer than 30 minutes (i.e., no click events originate from this person during that time), we simply say that the user has exited the site.

This assumption allows us to group these click events from the original datasets into user sessions, which illustrate the path a user takes while browsing the website and can be used to identify areas that attract more (or less) traffic.

Figure 2 illustrates one individual session chosen at random from our dataset. We can see that the user in this case was referred to our domain through a link that he or she found on a social network website and that their visit consisted of several hops, most of which happened in the *news* section.

Aggregating these sessions allows us to visualize which areas of the website are more popular, as well as which links connecting different sections are traversed the most. Take for instance the example illustrated in Figure 3. To generate this particular graph, we isolated the sessions corresponding to a certain newspaper's website, its 12 most popular sections, and the traffic between them. Among other observations, we noticed that the readers of this particular newspaper were often prone to navigating to the *sports* section and reading multiple articles there.

Furthermore, these sessions can be aggregated, producing a high-level view of the entire website structure by popularity of section. Figure 3 illustrates this concept.

Lastly, we note that based on information retrieved from specific features of our dataset, it is possible to determine if a user is simply browsing text articles, displaying image galleries or streaming online video. The following sections of this paper will describe how we used this fact to aid in the development of predictive models for video viewership engagement.

## METHODS

### Identification of Video Exit Instances

When a user watches a video, a separate log entry is made corresponding to when he or she completes watching a certain percentage of the video, while a player ID remains constant. This makes the clickstream log reflect a cumulative history of the viewer's progress within that video.

By filtering the data to get only clicks corresponding to video instances, and then by IP address, we obtain the entire video viewing activity of each IP. From this modified dataset, we isolate an individual "video view" table by specifying the player ID. This table is then sorted chronologically and filtered by session timeout. The last entry corresponds to the viewer's exit point. This gives us a unique session for a visit. In combination with the current session data, and data from cookies, we retrieve the user's unique historical browsing patterns. It should be noted that in the absence of cookies, we treat the user as a fresh incoming visitor. For our analysis, we isolated only the instances where the user exited the video. Due to the inherently discretized nature of the data collection, we get a coarse-grained estimate of when the user reached a certain percent of the video. If the last entry shows that the user watched 50% of a video, it can be inferred that the user exited at  $p\%$ , such that  $p \in [50, 75)$ .

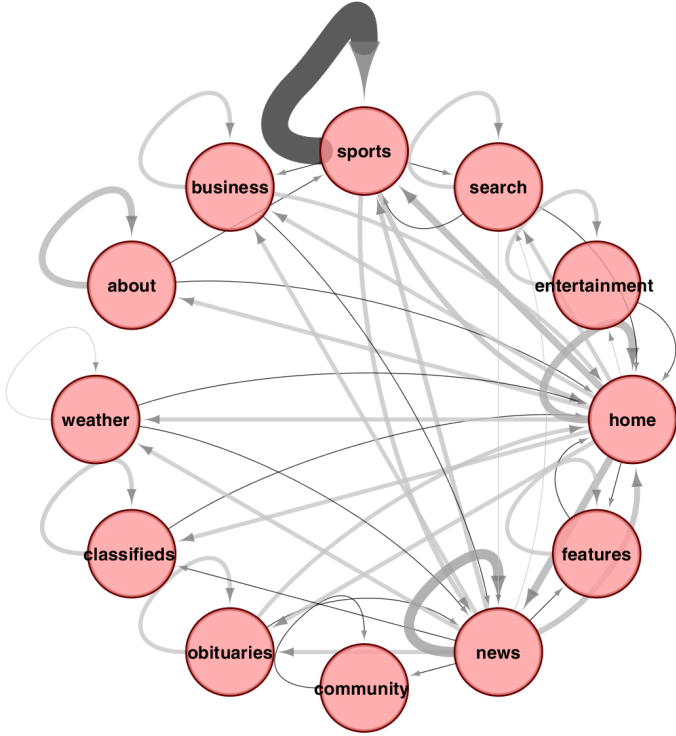


Figure 3. Clickstream network for a news-media website. The various nodes displayed here represent different sections. The direction of the arrows represents user traffic flowing between these sections and the thickness is indicative of volume of said traffic.

### Feature Selection

Using various feature selection methods, we reduced the size of our dataset from the original 161 features to the 12 best descriptors. Among these were features like IP, location, content annotations, and referrer information. Out of the 161 features in a typical video exit instance, 40 are mutually redundant, and 32 are constant in value. This motivates the need to find a set of features that is the best descriptor of the target class (in this case, the percent of video the user watches before exiting) [14]. We investigated various feature selection methods which support mixed data types and ranked the top features. One would expect these features to encompass measurable user traits which influence their interest in the video.

Various feature selection methods aim to remove redundant and irrelevant features using different statistical means, which have their respective strengths. Though a popular choice in machine learning, correlation based feature selection (CFS) was not considered due to the sparse nature of the data [15]. A more detailed study of these can be found in [28, 12]. The feature selection methods employed in this problem are described below:

#### Chi Squared

The chi squared ( $\chi^2$ ) method measures how much deviation is seen in the observed data from the case where the class value and the feature are independent of each other. It evaluates whether the feature and class occurrences are randomly related, or exhibit some relation.

Features	Chi	IG	GR	oneR	Symm
Time	1	1	7	-	2
IP	2	2	9	-	3
First hit referrer	3	3	5	2	5
First hit page	4	5	10	-	7
Story title	5	4	2	1	1
Search engine	6	7	3	3	8
City	7	6	-	-	9
ISP	8	8	-	-	10
Referrer type	9	10	1	-	4
# Pages viewed	10	9	8	-	6
Search page num	-	-	4	4	-
Frequency of visits	-	-	6	5	-

Table 2. Feature Selection Rankings. (description of abbreviations)

#### Information Gain

Information gain [23] measures how much entropy is lost when the feature is present vs. absent.

#### Gain Ratio

Information gain favors attributes with many values over those with fewer values, the gain ratio [24] compensates for this by factoring in the amount of split caused by the feature.

#### One R

One R formulates a set of simple relationships between the features and ranks the features based on how accurate these rules are.

#### Symmetric Uncertainty

Symmetric uncertainty [26, 10] targets attributes which correlate well with the class but have little intercorrelation.

The results of these feature selection methods are summarized in Table 2. The attributes in the table are the ones which consistently appear in the top 10%. These are the attributes which influence video exit points the most.

The time of viewing influences at what point people are prone to exit the video. IP address, in conjunction with location, and ISP indicate who is watching the video and thus offer a personalized facet to the prediction. The number of pages viewed by a person and frequency of visits can be perceived to be reflective of the person's interest in the site. The referrer which brought the viewer on the site can influence the engagement of the viewer; a viewer coming from a social network link interacts differently than one who had the site bookmarked on their browser. The entry point is the first page the viewer saw in their current viewing session; this determines their interest in consuming further content. The actual title of the story includes the section which the video is under. As we had observed in Figure 1, users viewing "Technology" related videos were less likely to exit than those viewing "Entertainment" related videos.

#### Classification

Our aim is to predict how much of the video a user watches before exiting. In our dataset, we find that this is represented by 5 distinct markers, which correspond to the percentage of

the video the user watched before exiting. We formulate two classification tasks- to predict what percent of the video is watched, and whether the user exits the video “early” (before reaching 50% of the video).

This prediction task involving 5 classes. Since it is relevant to predict users who exit early on in the video, we assume that users who exit the video at the beginning or having viewed 25% of the video to have exited “early”. As described above, this would correspond to users who have viewed 0 to 49% of the video.

	Predicted Class				
	Video Start	25% Complete	50% Complete	75% Complete	Complete
Actual Class	Video Start	✓	✗	✗	✗
	25% Complete	✗	✓	✗	✗
	50% Complete	✗	✗	✓	✗
	75% Complete	✗	✗	✗	✓
	Complete	✗	✗	✗	✓

	Predicted Class				
	Video Start	25% Complete	50% Complete	75% Complete	Complete
Actual Class	Video Start	✓	✓	✗	✗
	25% Complete	✓	✓	✗	✗
	50% Complete	✗	✗	✓	✓
	75% Complete	✗	✗	✓	✓
	Complete	✗	✗	✓	✓

**Figure 4. Converting the Percentages Classification to Early Exit Classification: The 5 class problem (top) is reduced to a binary classification problem by merging classes (bottom).**

We can refine the problem as the binary prediction of these “early exits”. The classes would then be a merger of the previously mentioned 5 classes, with the first two combined to form that of “early exits” and the latter 3 being those who chose to not exit early. This simplification is depicted in the representative expected confusion matrices for both of the classification tasks as displayed in Figure 4. We have performed both prediction analyses on our data.

## Naive Bayes

Among the simplest and most primitive classification algorithms, this probabilistic method is based on the Bayes Theorem [2] and strong underlying independence assumptions. That is, each feature is assumed to contribute independently to the class outcome.

## C4.5 Decision trees

C4.5 Decision Trees [24] work by building a tree structure where split operations are performed on each node based on information gain values for each feature of the dataset and the respective class. At each level, the attribute with highest information gain is chosen as the basis for the split criterion.

## Repeated Incremental Pruning to Produce Error Reduction

RIPPER [5] is a rule based classification tree learner. It is algorithmically faster than C4.5, having a complexity of  $O(n(\log(n))^2)$  as opposed to C4.5’s complexity of the order  $O(n^3)$ . RIPPER constructs an initial set of rules and then iteratively optimizes it according to a tunable parameter. It is implemented in Weka under the “JRip class”.

## Random forests

Random forests [3] combine multiple tree predictors in an ensemble. New instances being classified are pushed down the trees, and each tree reports a classification. The “forest” then decides which label to assign to this new instance based on the aggregate number of votes given by the set of trees.

## Decision Tables

Decision Table classifiers [18] are built by concatenating a series of rules derived from the feature set to corresponding class outcomes. This method as its major advantages the fact that it is easy to interpret and notably efficient.

## Random Subspaces

The random subspace method [17] is an ensemble classifier whose individual classifiers operate on random subsets of the feature set. The predictions made by the individual classifiers are combined using the posterior probabilities of each class in the constituent classifiers. This method looks at the classification problem from various perspectives by randomizing the selection of features.

## Stacking

Stacking [27] is a meta-classification scheme which employs an ensemble of classifiers and performs the learning task on two levels. First, the classifiers in the ensemble are trained on the data, then the meta-classifier learns from *their* predictions and the training labels of the data.

## Key Performance Indices / Metrics Utilized

Our key performance index is the accuracy of prediction of when the user will drop-off in the video. To obtain these predictions, we perform 10-fold cross-validation on the available data using various classification methods. In 10-fold cross validation, the data is randomly partitioned into 10 subsets and predictions are made on each of these. These predictions are then aggregated to provide the overall performance of the classifier, which we measured by the accuracy and area under the Receiver Operating Characteristics curve (AUROC), all



Dataset	Classifier	Acc	AUROC
Multiclass	NB	0.416	<b>0.718</b>
	C4.5	0.547	0.699
	RIP	0.547	0.629
	DT	0.543	0.717
	ST	<b>0.569</b>	0.652
Binary	NB	0.772	0.753
	C4.5	0.809	0.794
	RIP	0.826	0.697
	DT	0.806	<b>0.805</b>
	ST	<b>0.846</b>	0.739

**Table 3. Summary of results obtained for each classifier and dataset. The classifiers used are NB: Naive Bayes, C4.5: C4.5 decision tree, RIP: Repeated Incremental Pruning to Produce Error Reduction, DT: Decision Table, ST: Stacking using random subspaces of decision trees**

of which are described in further detail below. Each of these measures depict various aspects of the prediction results.

#### Accuracy

The accuracy of a classifier is perhaps the simplest measurement of its performance. It represents the percentage of total instances that were correctly classified. We would like for this to be as high as possible. The baseline for accuracy is that of a perfectly random prediction. For a binary classification problem, this would be 50% and for a 5 class problem, the baseline accuracy would be 20%. Any classifier which delivers statistically greater accuracy than these respective baselines, is considered to be better than a random predictor.

#### Receiver Operating Characteristics (ROC) curves

A system tuned to increase accuracy does not necessarily make it a good predictor. Relying on accuracy alone does not provide insights into the nature of misclassified instances. ROC curves [11] are a way to quickly compare multiple classifiers. The goal of a classifier in ROC space is to be as close to the upper-left corner as possible. In ROC space, if the curve for one classifier is closer to the upper-left corner than that for another, then it is considered to have a superior performance.

## EXPERIMENTAL RESULTS

We evaluated the performance of each of the classifiers used, with 10-fold cross validation for both the multiclass and the binary classification predictions. Table 3 summarizes the results of all experiments.

#### Multiclass Prediction


We see that in terms of sheer accuracy, the stacked classifiers performed slightly better than other methods, achieving an accuracy of 56.9%. In terms of AUROC, however, it is seen that Naive Bayes performs much better, closely followed by Decision Tables. These simple classifiers might not have the best accuracy, but outperform the others.

#### Binary Class Prediction

In this second scenario, we associate a semantic meaning to the drop-off percentage point and predict if the user will exit early or not. This refinement of the problem statement gives us a much better performance across the board. The stacked classifiers, for instance, achieve a remarkable accuracy of 84.6% when predicting which users exited their video streams prematurely. As it was the case with the multiclass problem, we again saw that Decision Tables and Naive Bayes surpassed the other classifiers in terms of AUROC values.

Though stacked classifiers give greater accuracy, they are not as good as Decision Tables or Naive Bayes in predicting early drop-off. This is still reflective of the general trends observed in the multiclass problem as we have merely merged classes, the underlying data remains the same.

In both, the multiclass and binary class prediction, it is observed that simpler rule based learners outperform complicated meta-classifiers. This is documented in [9], showing that stacking does not always outperform the best classifier.



**Figure 5. ROC curves for the binary class problem. A comparison of various classifiers to predict early exit behavior.**

We see that simple classification algorithms can be used to achieve comparable, or even better performance than complicated meta-classifiers. Besides the obvious performance superiority, it is desirable to use simpler classifiers on grounds of computational complexity, as implementing these is algorithmically more scalable and thus offers faster runtime.

## CONCLUSIONS

We demonstrated how clickstream data can be used to predict “early exits” in online videos. By constructing models to this effect, we were able to identify with high accuracy which video streaming sessions are likely to terminate prematurely. Additionally, we compared and contrasted the performance of a number of classifiers, highlighting those that we found to be

particularly fit to this problem. Having knowledge of such information would allow content providers to personalize how their media is distributed so as to increase user retention, and as a result, business value.

## REFERENCES

1. Banerjee, A., and Ghosh, J. Clickstream clustering using weighted longest common subsequences. In *Proceedings of the web mining workshop at the 1st SIAM conference on data mining*, vol. 143, Citeseer (2001), 144.
2. Bayes, M., and Price, M. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, communicated by Mr. Price, in a letter to John Canton, M.A. and F.R.S. *Philosophical Transactions (1683-1775)* (1763), 370–418.
3. Breiman, L. Random forests. *Machine learning* 45, 1 (2001), 5–32.
4. Bucklin, R. E., and Sismeyro, C. A model of web site browsing behavior estimated on clickstream data. *Journal of Marketing Research* (2003), 249–267.
5. Cohen, W. W. Fast effective rule induction. In *ICML*, vol. 95 (1995), 115–123.
6. Das, A. S., Datar, M., Garg, A., and Rajaram, S. Google news personalization: scalable online collaborative filtering. In *Proceedings of the 16th international conference on World Wide Web*, ACM (2007), 271–280.
7. Dobrian, F., Sekar, V., Awan, A., Stoica, I., Joseph, D. A., Ganjam, A., Zhan, J., and Zhang, H. Understanding the impact of video quality on user engagement. *SIGCOMM-Computer Communication Review* 41, 4 (2011), 362.
8. Dreze, X., and Hussherr, F.-X. Internet advertising: Is anybody watching? *Journal of interactive marketing* 17, 4 (2003), 8–23.
9. Džeroski, S., and Ženko, B. Is combining classifiers with stacking better than selecting the best one? *Machine learning* 54, 3 (2004), 255–273.
10. Eom, J.-H., and Zhang, B.-T. Machine learning-based text mining for biomedical information analysis. *Genomics & Informatics* 2, 2 (2004), 99–106.
11. Fawcett, T. An introduction to roc analysis. *Pattern recognition letters* 27, 8 (2006), 861–874.
12. Forman, G. An Extensive Empirical Study of Feature Selection Metrics for Text Classification. *The Journal of Machine Learning Research* 3 (2003), 1289–1305.
13. Guo, J. K. P. J., Krzysztof, D. T. S. P. M., and Miller, Z. G. R. C. Understanding in-video dropouts and interaction peaks in online lecture videos.
14. Guyon, I., and Elisseeff, A. An introduction to variable and feature selection. *The Journal of Machine Learning Research* 3 (2003), 1157–1182.
15. Hall, M. A. *Correlation-based Feature Selection for Machine Learning*. PhD thesis, The University of Waikato, 1999.
16. Hauser, J. R., Urban, G. L., Liberali, G., and Braun, M. Website morphing. *Marketing Science* 28, 2 (2009), 202–223.
17. Ho, T. K. The random subspace method for constructing decision forests. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 20, 8 (1998), 832–844.
18. Kohavi, R. The power of decision tables. In *Machine Learning: ECML-95*. Springer, 1995, 174–189.
19. Liu, J., Dolan, P., and Pedersen, E. R. Personalized news recommendation based on click behavior. In *Proceedings of the 15th international conference on Intelligent user interfaces*, ACM (2010), 31–40.
20. Mobasher, B., Cooley, R., and Srivastava, J. Automatic Personalization Based on Web Usage Mining. *Communications of the ACM* 43, 8 (2000), 142–151.
21. Moe, W. W. Buying, searching, or browsing: Differentiating between online shoppers using in-store navigational clickstream. *Journal of Consumer Psychology* 13, 1 (2003), 29–39.
22. Montgomery, A. L., Li, S., Srinivasan, K., and Liechty, J. C. Modeling online browsing and path analysis using clickstream data. *Marketing Science* 23, 4 (2004), 579–595.
23. Quinlan, J. R. Induction of decision trees. *Machine Learning* 1, 1 (1986), 81–106.
24. Quinlan, J. R. *C4.5: Programs for Machine Learning*, vol. 1. Morgan Kaufmann, 1993.
25. Srivastava, J., Cooley, R., Deshpande, M., and Tan, P.-N. Web usage mining: Discovery and applications of usage patterns from web data. *ACM SIGKDD Explorations Newsletter* 1, 2 (2000), 12–23.
26. Witten, I. H., and Frank, E. *Data Mining: Practical Machine Learning Tools and Techniques*. Morgan Kaufmann, 2005.
27. Wolpert, D. H. Stacked generalization. *Neural networks* 5, 2 (1992), 241–259.
28. Yang, Y., and Pedersen, J. O. A comparative study on feature selection in text categorization. In *ICML*, vol. 97 (1997), 412–420.